



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.


In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed 

Dated 14 August 2003



1/77

Request for grant of a patent
(See notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office
Cardiff Road
Newport
Gwent NP10 8QQ

1. Your reference

PHGB 020195

19NOV02 E764225-1 003008
P01/7700 0.00-0226803.5

2. Patent application number
(The Patent Office will fill in this part)

0226803.5

16 NOV 2002

3. Full name, address and postcode of the or of each applicant (underline all surnames)

KONINKLIJKE PHILIPS ELECTRONICS N.V.
GROENEWOUDSEWEG 1
5621 BA EINDHOVEN
THE NETHERLANDS
07419294001

Patents ADP Number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

THE NETHERLANDS

4. Title of the invention

STATE MACHINE MODELLING

5. Name of your agent (if you have one)
"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

Philips Intellectual Property and Standards
Cross Oak Lane
Redhill
Surrey
RH1 5HA
08359655001

Patents ADP number (if you know it)

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority Application number
(if you know it)

Date of filing
(day/month/year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day/month/year)

3. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer "Yes" if:

YES

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document.

Continuation sheets of this form

Description	20
Claims(s)	5
Abstract	1
Drawings	13

only 13

10. If you are also filing any of the following, state how many against each item:

Priority Documents

Translations of priority documents

Statement of inventorship and right
to grant of a patent (*Patents Form 7/77*)

Request for preliminary examination and
search (*Patents Form 9/77*)

Request for substantive examination
(*Patents Form 10/77*)

Any other documents
(*Please specify*)

11.

I/We request the grant of a patent on the basis of this application.

Signature

A. G. White

Date 14/11/2002

12. Name and daytime telephone number of
person to contact in the United Kingdom

01293 815438

(A. G. WHITE)

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.*
- Write your answers in capital letters using black ink or you may type them.*
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.*
- If you have answered "Yes" Patents Form 7/77 will need to be filed.*
- Once you have filled in the form you must remember to sign and date it.*
- For details of the fee and ways to pay please contact the Patent Office*

Patents Form 1/77

DESCRIPTION

STATE MACHINE MODELLING

5 This invention relates to a method of modelling a state machine comprising a first state model, and a second state model implanting a function call. The invention relates also to a computer program for instructing a computer to carry out the method, and to a computer programmed with such a computer program.

10

 The increasing complexity, size and lead time of software systems is a concern to the software industry. One solution which addresses these concerns is the component-based synthesis of systems. Components provide for system modularity, customisability, maintainability and
15 upgradeability. Large systems can be built by selecting components from repositories and binding them together, rather than by writing systems without code reuse.

 This approach to system building presents problems in terms of the testing of the operation of the system, although the testing of the individual
20 components is relatively straightforward. Currently available testing tools rely on a tester generating a state machine model of the system under test. The model and the system under test are then subjected to a test, and the resultant state of the model and the system compared to determine if the system performed as expected. This procedure is repeated until the system
25 has been fully tested.

 An explanation of state machines and how they are modelled now follows. It will be appreciated that the notation and the syntax used is illustrative and non-limiting. The following also describes how models are used for system testing.

30 Many systems can be modelled according to their state behaviour, that is their state and how the state changes as a result of some stimulus or signal, called an event. Under this modelling technique, if a system is in a

given state, it will remain so indefinitely until an event occurs. The notion of a state therefore entails durability – the state exists over a period of time. Even if a system enters a particular state s_1 and there is an event ready and waiting to cause a change of state (to state s_2), the moment when the system is in state s_1 is a point at which the system is stable in terms of its state behaviour? At such a point, the state of the system (in a wide sense) will map to a state in the model of the system.

Events are modelled as instantaneous signals which have no duration. They are able to trigger some processing in the system which may or may not result in a new state. In some states, events may be ignored by the system, leaving the system in its existing state.

A system may be of the kind that theoretically runs indefinitely, such as an operating system or real-time kernel, or it may have a clear lifecycle. However, even operating systems can generally be closed down in a controlled way.

A multi-threaded application might be modelled with states which represent the fact that low priority threads are running. Such a system would still be able to react to events which interrupt at a higher priority. It may even be necessary to represent cpu (central processing unit) bound tasks as states, perhaps using several states so as to model events as having been recorded but unable to be processed until the task completes.

Input data to a program can also often conveniently be thought of as a sequence of events. In this case, the program will normally have instant access to the next event (apart from an occasional disc-access or similar), and so will be cpu-bound, but this does not detract from the state model. An example of such a kind of program is a conversion program to convert texts from one kind of character coding to another, perhaps with situations where one character of input maps to two characters of output and vice versa. The input characters (including new lines and end-of-file) can be modelled as events. Output characters will be generated on certain state changes. Another example is a compiler where the input tokens can be

regarded as events; the state is some record of completed successful parsing of 'terms' in production rules.

If two states show identical responses to any sequence of events that is processed from a system in such a state, then they are indistinguishable and are best modelled as one state, so as to avoid redundancy in the model.

In order to model a system, it is necessary to express all the relationships between states, events, and new states after processing the event. A transition maps a source state to a new state (the target or destination state) by means of an event. In effect, the event triggers the transition. There can be multiple target states, but this is not discussed further here. A diagram showing states and transitions is termed a state-transition diagram. States are conventionally denoted by circles, and transitions by arcs with an arrowhead. Transition arcs conventionally are annotated with the events that cause the transition. Figure 1 shows a system having three states: a, b and c; four events: α , β , γ , and δ ; and four transitions: t1, t2, t3 and t4.

At any one time, a system modelled by the Figure 1 state-transition diagram will be in only one state. That state is called the occupied (or active) state. The others are vacant (or inactive). Transitions whose source states are vacant at the time an event occurs do not cause any state transitioning to take place – they are inapplicable in the current state. If an event occurs which is the trigger to a transition whose source state is occupied, then (apart from non-deterministic situations) the transition takes place. Here, the source state becomes vacated and the target state becomes occupied. In the Figure 1 example, when the system is in state a, it reacts to event α by executing transition t1, i.e. by transitioning from state a to state b. If the system is not in state a, then transition t1 is not applicable because the system is not in t1's source state. Only one transition takes place as a result of one occurrence of this event, so transition t2 does not take place as well, unless and until another event α occurs. It will be appreciated that there can be several transitions emanating from any state

(for example t_1 and t_3 from state a). Also, an event can be a trigger to more than one transition (for example α triggers t_1 and t_2), but, (excluding non-determinism), it is not usual to find two transitions triggered by the same event from the same source state. Furthermore, a transition can be triggered by more than one event, in which case any one of the events will trigger the transition. For example, transition t_3 is triggered by event β or δ . If an event occurs which does not trigger a transition, (for example if in state b event β occurs), then the event is disregarded and no state change occurs.

The way in which the state transition diagram of Figure 1 is represented in a possible source code language is:

```
statechart sc(s)
  event alpha,beta,gamma,delta;
  cluster s(a,b,c)
    state a {alpha->b;beta,delta->c;}
    state b {alpha->c;}
    state c {beta,gamma->a;}
```

Here, the state transition diagram is declared as a "statechart", which consists of a cluster s , which consists of three states or leafstates a , b , and c . A cluster indicates a grouping in which no more than one member state can be occupied. Events are declared and are used in transitions, which are denoted by

events -> target state;

State behaviour modelling is part of the UML (Unified Modelling Language) dynamic view.

An implementation of a state based model of a system provides a way to automatically generate test cases for that system. According to a white-box technique, this can be set up is for a test script (TS) to communicate with the State Behaviour Model (SBM) and the Implementation Under Test (IUT), giving each instructions to put themselves in a particular state, to

process an event, and to provide their new state. The test script then compares the new states as reported by each. Any mismatch is a test failure and so possibly the detection of a bug in the IUT (although it could be a modelling error, a test script error, or even a bug in the SBM). The instruction sequence is repeated for as many states and events as it is feasible to execute. A certain amount of glue-code is needed to communicate with the IUT.

The SBM resultant state is termed the 'oracle' to the test, i.e. it is the expected result from the IUT. The process is illustrated in Figures 2A and 2B.

This technique is called white-box because it requires knowledge of the internals of the IUT in order to communicate with it in this way. Black-box techniques also exist in which the IUT is entirely event driven and its behaviour is deduced from limited observed output traces which are generated on certain transitions. In this case, transition tour algorithms provide some form of coverage of the state space.

Coverage of all states and events is obtained by looping as follows:

For all states

For all events

Set state in SBM and IUT


Process event in SBM and IUT

Get state of SBM and IUT

Compare resultant states




This does not guarantee the correct state behaviour of the IUT - it is possible that it could show incorrect behaviour in some states under some circumstances - e.g. entering a state via one route might give rise to different subsequent state behaviour than the state behaviour when the same state is reached via a different route. Other test generation algorithms can be designed to give further coverage, for example covering all pairwise event combinations.

Tests are preferably called in a uniform way, and each test should provide its own pass/fail criterion. The test report should produce a uniform description of whether each test passed or failed. A tool providing facilities for doing this is called a test harness.

5 A more detailed explanation of state machines follows. The example state machine of Figure 1 contains three leafstates which are "wrapped" in a cluster. A cluster is a group of states (members of the cluster) such that at most one member state can be occupied. If one member is occupied, the cluster is regarded as occupied. If all members are vacant, the cluster is
10 vacant. The members of a cluster can be other clusters, sets or leafstates. The diagrammatic notation for a cluster is a rounded rectangle with its name at the top left. One member of the cluster is designated the default member (symbol ). If and when the cluster is initially entered or is the target state of a transition then, unless other factors come into play, the default
15 state is entered.

Transitions can have a cluster as their source state. They can also have a cluster as a target state. This gives a compact way to express what would otherwise be multiple transitions. An example cluster is shown in Figure 3A. The equivalent flattened state machine is shown in Figure 3B.

20 A cluster can be marked with a history or deep history marker. The history data records the member that was occupied when the cluster was last occupied. On diagrams, history is marked according to the following legend:

 = no history (default)	 = (shallow) history	 = deep history
--	---	--

25

A cluster with a history marker, when entered without a specific member being specified, will enter the historical state. If history data is not available, the default state will be taken. Deep history indicates that historical data is to be used (assuming it is available) on re-entering the cluster and all

descendant clusters below the marked cluster. The descendant clusters are entered under a deep history obligation – whether or not they have a history marker. The deep history obligation is not applicable simply because a particular cluster is below another one with a deep history marker. It must be the case that the cluster with the deep history marker is actually entered in the course of the transition for the deep history obligation to apply. History data can be cleared by a function call.

A set is another means by which states can be grouped hierarchically. If a set is occupied, all its members must be occupied. If the set is vacant, all its members must be vacant. The members of a set can be clusters, sets or leafstates. A set normally has at least two members, which provides a statechart with concurrency (i.e. parallelism): several states can be occupied in parallel. The notation for a set is a rounded rectangle with a tab. Members of a set are separated by a dotted line. A separate enclosing rectangle around the cluster is not required; the symbol in the member area but not in any other symbol indicates a cluster. Figure 4 shows how members of sets can be designated.

A set cannot be marked with a history marker, since there is no choice as to which member to enter – if a set is entered, all of its members are entered. A set can be marked with a deep history marker. This means that on entry into the set and then into the set members, a deep history obligation will be passed on to all members of the set. Any clusters below the set in the hierarchy will then be entered in their historical state.

Other features in state machines are listed below:

Conditions on transitions. For example, ' $\alpha[V1 \geq V2]$ ' on a transition means that the transition will only be made on event α if V1 is greater than or equal to V2. The conditions may be C-like boolean expressions, and may contain tests for the occupancy/vacancy of other states.

Actions on transitions, including firing additional events, executing embedded imperative code (such as C code). For example, ' $\alpha/\text{fire_f1}$ ' on a transition means that event α will trigger the transition, and the transition will fire event f1.

Lambda-events, i.e. events that are generated automatically (typically by forward chaining) when some condition becomes true such as when some variables take on some particular values.

5 **Meta-events**, i.e. events which take place when some micro-step in transition processing takes place, such as entering or exiting some particular state.

Use of **scoping operators**. It may be desirable to allow e.g. state, event, and variable (and, for a typed language, tag-) and other names to be identical, but to be scoped to different parts of the state hierarchy. A rough
10 similarity comparison can be made with global and class member variables in C++, where the :: operator accesses the global variable, but in a statechart precision is required down at any hierarchical level. Scoping operators give access to the desired target name. Possible scoping operators are

15 • **descend** (diadic, right-associative, infix, higher precedence), e.g. a.b.c means descend through state a down through state b to state c.

• \$ **back** (monadic, right-associative, prefix, lower precedence). Back out one level and address a state from that level. \$\$a
20 means back out two levels and enter state a.

• :: **fromtop** (monadic, right-associative, prefix, lower precedence). Back out to the outermost hierarchical level and address a state from that level. ::a.b means enter state a, then from there state b, from the outermost hierarchical level.

25 • %% **ancestor** (diadic, right-associative, infix, higher precedence). Back out to a named level (the left-hand operand), then enter the state(s) denoted by the right hand operand.

The above operators typically have precedences higher than those of arithmetic operators. Additional operators will be known by those skilled in
30 the art, including operators taking an implicit this-state argument.

Currently available tools for producing state machine models are unable to model function calls accurately. A software component is accessible only

through its interfaces, each interface being a collection of function calls. An operating system is an example of a software component. A function call can be considered to be made from a client to a server, a server being one who provides a service. Function calls can be synchronous or asynchronous. Binding between components A and B, one of which is a client and the other a server is shown in Figure 5.

Supposing a model of the state behaviour of component A and of component B exists, a problem subsists in how to obtain a model of the state behaviour of the combination of the two. This model would involve name identification (in the sense of making the names the same) between caller and callee, which can be done in the prior art. A conventional way to model function calls in a finite state machine is for each function to be a parallel machine, so that each function is available once. An example of this is shown in Figure 6. In Figure 6, the processing of the maximum function is artificially regarded as going through various intermediate states, and as making various additional calls (to f1 and f2). In this example, the server calculates the largest (maximum) value of the variables P1, P2 and P3, and returns the determined value to the client, by firing event `ret_maximum`.

Supposing f1 were implemented on the client side above, f1 clearly would be best modelled by a parallel machine (i.e. set member) on the client side. Furthermore, if f1 were to call a function g1 that was implemented back on the server side, then f2 as well would be best modelled by a parallel machine on the server side. It appears to be a general rule that all functions should be modelled as parallel machines within their component wrapper state. This is shown in Figure 7.

In practice, the functions f1, f2, f3, f4 of Figure 7 typically would have different signatures and state behaviour. The figure shows the general structure of a main component state and functions modelled as co-members of the set.

Functions which return a pending status and notify later may need to be modelled in a way that accepts any order of notifications. In the example of

Figure 8, the client makes two calls on a thread which cannot be interrupted on the client side. As soon as more than two functions can be outstanding, it becomes impracticable to extend this concept. Combinatorial explosion takes place of the number of states representing some, but not all notifications. A more general scheme is shown in Figure 9.

There are limitations to the above described function call modelling schemes, in that they cannot handle re-entrant and recursive calls. Re-entrancy occurs when a first call to a function is incomplete and a second client makes a call to that function, typically on its own thread. Recursion occurs when a function directly or indirectly calls itself. It is an aim of the invention to address these shortcomings.

According to the present invention, there is provided a method of modelling a state machine comprising a first state model, and a second state model implementing a function call, the method comprising, in response to an event in the first state model instructing the firing of the function call, implanting the second state model in the first state model.

According to a second aspect of the invention, there is provided apparatus for modelling a state machine comprising a first state model and a second state model implementing a function call, the apparatus comprising means responsive to an event in the first state model instructing the firing of the function call, for implanting the second state model in the first state model.

Using the invention, it is possible to model accurately function calls without being susceptible to combinatorial explosion of the number of states, whilst being able to deal with re-entrant and recursive calls.

In the embodiments, a function call is modelled in a second state machine which is independent of the first state machine. When the first state machine calls the function call, for example using a leafstate "calling", the second state machine is temporarily implanted over the "calling" state. Static recursion or infinite compile time recursion is avoided since the implantation is made only at the time of calling the function call, rather than

at compile time. After entering the state machine, return is made to the first state machine after transition to a terminator state, which fires an event called \$return (where \$ indicates scoping back one level, and 'return' indicates an event which fires the transition to state "after"). This can be described as a synchronous function call. In an asynchronous function call embodiment, a second state model is implanted in free-space, and has a lifetime which is independent of any other state machine model. An asynchronous implanted state machine returns a notification upon completion, i.e. on transition to a terminator state. Intermediate notifications may also be given. On exiting any implanted function call state machine, the second state machine is deleted.

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, of which:

Figure 1 is an example state diagram;

Figures 2A and 2B show how a test script tests on implementation under test (Figure 2B is a message sequence diagram);

Figures 3A and 3B illustrate the effect of a cluster;

Figure 4 illustrates set and set member notation;

Figure 5 illustrates how components can be bound;

Figures 6, 7, 8 and 9 illustrate how function calls can be modelled in the prior art;

Figures 10 and 11 show function call modelling according to the invention;

Figures 12 and 13 show configurator modelling according to the invention;

Figure 14 shows a bound function call modelled according to the invention;

Figures 15 and 16 show unbound function calls modelled according to the invention;

Figures 17 and 18 illustrate the use of synchronous function call modelling according to the invention;

Figure 19 shows how a chain of transitions may be created using synchronous function calls according to the invention;

Figure 20 illustrates the use of asynchronous function calls modelled according to the invention; and

5 Figure 21 illustrates how the model according to the invention may deal with both synchronous and asynchronous function calls.

The invention is concerned with the modelling of state machines. The model is formed from program code prepared in any suitable programming
10 language and running on any suitable machine. For example, the model may be programmed in Prolog™ and be run on a general purpose computer. The invention resides in the modelling, and it will be apparent to the person skilled in program how to generate suitable program to implement the invention.

15 Figure 10 shows a state machine model according to the invention. A function (f1) is modelled as a cluster named mf1, and is free-standing, i.e. it is not attached to any part of the statechart hierarchy. It is marked as recursive, though this might be implicit for free-standing members. A 'calling' event (mf1.ef1) is recognized as being special, because something
20 special happens when its transition takes place.

The special effect is as follows:

- The target state ('calling') is temporarily overwritten (until a 'return' event) by an instantiation of the cluster mf1 *in situ*. This could be termed 'implanting' a new part of the state machine on top of another state.
- 25 - the transition triggered is processed in the instantiated member mf1. When the event can trigger several transitions within the instantiated member (perhaps with different event parameter signatures, or different conditions on the transition), each transition is dealt with as if the new member were a permanent fixture of the state machine.
- 30 - When the instantiated member fires a '\$return' event (which is notation for a return event in the parent class's scope), the instantiated member mf1

is removed, leaving the original 'calling' state standing. This models the function return.

It will be noted that: The client declares a local event called 'return', to be used for all function returns. The symbol \nearrow represents an event declaration which the function member mf1 scopes precisely by firing
5 'return' (the \$ indicating parent scope). Thus, even with many function calls active, possibly from several clients and at various levels of nesting, there could never be any ambiguity concerning which 'return' is being referred to.

Also, the function-call event is modelled here as local to the member
10 being instantiated, although it could instead be modelled in another way. The client directs the event to the member using a 'descend' operator (the dot), giving mf1.ef1..

Furthermore, once the function member mf1 has been instantiated, it is just like a normal fixture of a state machine model.

15 The special nature of function invocation could be recognized by the fact that the fired event (mf1.ef1) scopes to an event in a recursive member (mf1) or by the fact that the fired event is an action on a (client) transition which targets a state named 'calling'. Such a state could be considered a marker state.

20 A more complex example is shown in Figure 11. In this figure, client3 calls a function f3, which calls a function f4, which calls the function f3, at which point a snapshot is taken. The mutual recursion in this model can be broken in the function f4, which need not call f3. Since a function member is instantiated only at the time of the transition to it, not at compile time, no
25 infinite recursion occurs in this example.

It will be appreciated that this allows proper modelling of immediately recursive calls, indirectly recursive calls, deep nesting, and calls from several clients.

30 The above described technique in effect uses infinite state machines, rather than the conventional finite state machines, since it is theoretically capable of containing arbitrarily many states.

To model the action on the configurator of Figure 5, including its creation of the two components, it is possible to allow the state machine model to modify itself to represent the new state structure on creation of the components. This is illustrated in Figures 12 and 13. In Figure 12, a transition from a leafstate j1 to another leafstate j2 in a cluster jmain causes code to be run, which code causes the creation of set k. The state machine model resulting from the code running event triggered by the transition is shown in Figure 13. As an optimisation, the dynamic 'compile' action could incorporate pre-compiled program code.

In summary, a state machine model according to the invention allows for recursive set members. These have the property of always responding to a particular starting event, cloning a new statechart member every time a call is made. Recursive set members could be denoted by the symbol



If there are several instances of a called function, the state machine execution logic is required to direct certain events to the right recipient or recipients. Calls from separate clients behave similarly to calls on separate threads, each with its own stack. Calls to 'from function to function' and their 'returns' are like stack push-and-pop operations on the same thread. Return events are required to be directed to the deepest caller on any particular thread. This is very probably directly analogous to what is happening in the system which is being modelled. Global events are required to be recognized as such since they are distinct from return events and do not entail any analogy to a push or pop operation. Global events can trigger transitions indiscriminately.

The concept of recursive set members can be applied to synchronous and to asynchronous situations. Synchronous function calls are considered here first.

In a synchronous function call, the function call executes on the caller's thread. When the function call returns, it is regarded as complete. Clearly, the caller cannot make a second function call until a synchronous call

completes, since the thread of control is not available. Within the category of synchronous function calls, it is possible to distinguish between bound calls and unbound calls.

5 A bound function call runs to completion without requiring any more events to drive it to completion. Such a call is typically CPU-bound – an example is a function to find the longest or maximum of a list of numbers. Alternatively, the bound function call might involve an activity which frees up the CPU (e.g. by performing some I/O (input/output)), but the execution is regarded in the model as predetermined rather than dependent on the
10 presence of an event. If the function call is not modelled as an atomic occurrence, i.e. if there can be an intervening event between start and completion, then the function call is better modelled as an unbound function call.

A bound function call may be modelled as a simple library function in an
15 assignment action on a transition as shown in Figure 14. Here, event α causes a transition to leafstate b although the transition can only be completed after the setting of a variable m to equal the maximum value of variables a, b and c.



A bound function call is capable of firing other events. A standard library
20 module can easily provide for this.



An unbound function call requires the occurrence of at least one additional event to drive the function to completion. For example, if a function obtains input from a user, it might be modelled as requiring the event "input_obtained" to complete it. Here, the case of restricting the user
25 to one function call per transition is considered initially. Single implantations are made per set member on the call. Figure 15 illustrates the process. In Figure 15 and the following figures, fire events on transitions are abbreviated to 'f', 'g', '\$return' and so on, i.e. the word 'fire' is omitted.

In Figure 15, the optional fired events on the right hand side marked
30 "optionally also ::f_return" and "optionally also ::g_return" are included for completeness because they would be useful in the event of a pumped call.

The notation "::" indicates global scope. Pumped calls need not be discussed here.

Two options are possible for representing the client side call –one with an intermediate "calling" state and one without (both are illustrated in Figure 16) but both corresponding to the same implementation in principle. The representation of Figure 15 is expected to be more developer friendly.

The arrow  in Figure 15 indicates the dynamic transformation of the state machine when the event α occurs. A new state machine element is implanted on the calling event - indicated by . It is not necessarily adequate to perform a static implantation of the function (e.g. at compile time) as some functions are recursive; it is only at run-time that the termination condition can be identified and executed. It will be noted that the tip of the transition is to the newly created element not the state indicated by the user, which is reached only on function completion. The clusters marked <S> are implanted on function call, and are removed on completion.

The standard notation  shows what state is entered on creation of the implantation. Function calls do not remember state history from one invocation to another. The notation  (called a terminator) indicates that the implantation is to be removed after a transition entering the terminator. There could be several different transitions to the terminator. There should only be one initial state in implantable machine models, just as only one default state is allowed in conventional state machines. If necessary, fork nondeterminism could be applied on the next transition. Although the implantations shown in the example are clusters, they could be sets, or even leafstates.

Whatever the option chosen for the user representation in the state machine model, the implementation may make use of the technique used by asynchronous calls. This is shown in Figure 17. In Figure 17, the user cannot make use of any apparent asynchronous functionality, because the state marked 'calling_g' is not accessible for any other transitions. The client is locked in the state 'calling_g' until the return event is fired within function call g.

If several synchronous calls are put on a transition, they are interpreted in the model as a sequence. The implementation translates this into a chain of transitions, either built up in one go or created step-by-step. The principle is illustrated in Figure 18, in which an event α fires functions f and g . The function calls f and g are instantiated in a chain between the originator state before $_fg$ and the target state after $_fg$.

If a sequence of actions contains a mixture of function calls and other actions, for example:

$$\alpha/f, x=p+q, g, y=z, \beta, h, x=x+1$$

where f , g , and h are function call events, but β is just a global event, then chaining puts the non-function-call actions on the linking transitions. This is shown in Figure 19. The sequence (micro-step) in which the non-function-call actions, including the initial one, take place, is dependent on the transition algorithm. Possible implementation techniques are to make all implantations in the chain in one go, or to make them at the latest possible time, i.e. just before they are needed.

An asynchronous function call, on the other hand, provides return of control to the caller, here termed a pending return, but the function continues to do some (or all) processing on another thread. When the processing associated with the function call is complete, the function call provides a notification. It is possible that intermediate and final notifications could be given, indicating the completion of certain phases of processing, but any final notification means that no more notifications can come from this function invocation. The intermediate notifications can be regarded as ordinary broadcast events. In this connection, broadcast events are the same as fired events.

An asynchronous function call may, depending on run-time circumstances, provide either a synchronous-like completion, or a pending 'return' with 'notification' later. A request for a web page is used here as an example. If the page is in cache, it may be returned quickly with completion. Otherwise, the function call returns 'pending', accesses the page over the Internet (for example), and 'notifies' when the page has been obtained.

With the model according to the invention, it is possible to have several asynchronous function calls outstanding at any one time, each function call running on a respective thread. For this reason, in a state machine model, the caller and all called functions are able to transition independently. Also, by parameterising the events, or by using names related to the function name, provided is a means to distinguish pending, notify and other events. This avoids ambiguity as to which function produced it.


When an event representing an asynchronous function call is processed, a state-machine element is implanted, but with a special status indicated here by a double perimetral line. The special status ensures that the implantation has a scope that is effectively local to the caller, and a lifetime that is independent of that of caller. For multiple function calls, a temporary implantation of a machine element is performed for every function called. An implementation of asynchronous function calls is illustrated in Figure 20. In this figure, a transition from leafstate 'before' to leafstate 'f&g' fires asynchronous function calls 'f' and 'g', which results in the state machine model shown on the right hand side of the figure. Here, function calls 'f' and 'g' are implanted or instantiated in free space but within the set named 'client'. In the function call 'f', an event 'pending_f' is fired as a leafstate 'fa' is entered, which causes a transition from leafstate 'f&g' to leafstate 'f_pnd'. In function call 'g', an event '::pending_g' is fired only on transition from leafstate 'gb' to leafstate 'gc'. This causes transition from leafstate 'f_pnd' to 'g_pnd'. An event '::final_notif_g' is fired on entering the terminator from either of leafstates 'gb' and 'gd'. In this Figure, behaviour is modelled via two routes: if f fires 'pending' before g then transition is made to state f_pend; if g fires pending before f does so, no transition is made either until both functions notify or until f fires 'pending'.

With asynchronous function call modelling, notification events are global (indicated by the "::" operator), which ensures that all recipients see them. The names of the notification events are required by the model to be unique. The optional fired events on the right hand side of Figure 20,

marked "optionally also \$async_return" are included for completeness because they are needed in the event of a pumped call.

Each of the implanted machine elements, marked **<A>**, has the scope of being a sibling of the element at the tip of the transition arrow. The scope is not related to the effective source leafstate, nor any orbital state. (Orbital states are states above the common ancestor of source and target states, where a transition specifies that states must be exited and re-entered up to this orbital state). This fixed policy facilitates precise targeting of broadcast events and variables when scoping operators are used.

The implanted machine elements, marked **<A>** and having a double perimetral line, can be regarded as extra active members of their parent. Thus, in the case of the parent being a cluster, the implanted machine elements break the ordinary rule that only one active member is allowed in a cluster. However, the implantation can be thought of as independent of its machine-path parent, since it has an independent life-cycle. The parent may not have the implantation marked as a child, so the implantation will not take part in algorithms which examine a parent's children. In this way, a cluster can be exited, for example, without interfering with the life-cycle of the implanted member.

The lifecycle of an implanted asynchronous function is independent of the lifetime of any other machine. It is destroyed only when the transition to the terminator (symbol ) takes place. Even if a parent is another function and is destroyed, the asynchronous function remains alive until it transitions to its own terminator.

Additional intermediate notifications can be included as well as a final notification. Only the final notification corresponds to the function implantation being removed; the modelling system knows that the implantation is to be removed from the fact that the transition is a transition to a terminator. From an implementation perspective, even a final notification behaves like any other broadcast event.

As mentioned above, events such as "pending" and "notify" need to identify the function they apply to by their name. Accordingly, events are

parameterized, which disambiguates between broadcast events from each of a function called twice (or more) in the same scope. (There are various ways in which some function can be called twice in the same scope).

5 More deeply nested parts of the implanted machine, if present, can use a \$\$ scoping operator, for example, when targeting fired events at the caller.

It will be noted that a generic transition on $\text{in}(\text{f_notif}) \ \&\& \ \text{in}(\text{g_notif})$ and other specific transition paths out of the calling state are possible.

10 If a transition contains calls to synchronous and asynchronous functions, implantations of the relevant kind can be provided systematically. In the example of Figure 21, events σ, τ, ν are synchronous functions, events α, β are asynchronous functions, and other events ($\epsilon, \zeta 1\text{--}\zeta 6$ and return events) are ordinary global events. Here, on transitioning from leafstate 'before' to leafstate 'after', a chain of synchronous function calls are
15 implanted between the two leafstates, and two asynchronous function calls are implanted in free-space. It will be seen that the function calls and events fired by the transition are handled by the model in the order in which they are listed. The transition between any two sequential synchronous function calls fires the events that are listed between the calling of the
20 functions on the pre-implantation model. The reaching of the leafstate 'after' is not dependent on either of the asynchronous function calls reaching their terminator and thus completing.

Here, the synchronous function calls may be implemented using the asynchronous function call modelling technique.

25 The invention may be carried out using any suitable computer, such as the personal computer of Figure 22. The computer 220 comprises a control processing unit 221, which runs a state machine modelling program stored on a hard disk drive 222. During the running of the program, a model of the state machine is built and stored in RAM 223.

CLAIMS

1. A method of modelling a state machine comprising a first state model, and a second state model implementing a function call, the method comprising, in response to an event in the first state model instructing the firing of the function call, implanting the second state model in the first state model.

2. A method according to claim 1, in which the second state model is absent of history information.

3. A method according to either preceding claim, in which the second state model contains one or more clusters.

4. A method according to any preceding claim, in which the second state model contains one or more sets.

5. A method as claimed in any preceding claim, in which the second state model contains two or more leafstates having one or more event driven transitions therebetween.

6. A method as claimed in claim 5, in which one or more of the transitions fires a notification event .

7. A method as claimed in any preceding claim, in which the second state model is implanted over an explicit marker state of the first state model.

8. A method as claimed in any of claims 1 to 6, in which the second state model is implanted over an implicit marker state of the first state model.

9. A method as claimed in claim 7 or claim 8, in which the second state model is deleted on completion.

5 10. A method as claimed in claim 9, in which the model allows the entering of a state in the first model local to the caller of the second state model only on deletion of the second state model.

10 11. A method as claimed in any of claims 7 to 10, in which local declarations and/or scoping operators are used in the second state model.

12. A method as claimed in claim 11, in which a return event from the second state model uses a "back" scoping operator.

15 13. A method as claimed in any of claims 1 to 5, in which the second state model is implanted in free-space.

14. A method as claimed in claim 13, in which the lifetime of the second state model is independent of any other model.

20 15. A method as claimed in claim 13 or claim 14, in which the second state model is implanted local to the caller of the second state model.

25 16. A method as claimed in any of claims 13 to 15, in which notification events from the second state model are global.

17. A method as claimed in any of claims 13 to 17, in which the second state model is deleted on transition to a terminator forming part thereof.

30 18. A method as claimed in any preceding claims, in which events occurring in the second state model are parameterised.

19. A method as claimed in any preceding claim, in which events occurring in the first state model are parameterised.

20. A computer program containing instructions for a computer to carry out the method of any of claims 1 to 18.

21. A computer program as claimed in claim 19, further comprising instructions for a computer to generate an executable program exhibiting the same behaviour as the state model.

22. A computer program as claimed in claim 19 or claim 20, further comprising instructions for a computer to generate tests with an oracle, for testing an implementation conformant to the behaviour of the state model.

23. A computer programmed with the computer program of claim 19.

24. Apparatus for modelling a state machine comprising a first state model and a second state model implementing a function call, the apparatus comprising means responsive to an event in the first state model instructing the firing of the function call, for implanting the second state model in the first state model.

25. Apparatus according to claim 24, in which the second state model contains one or more clusters.

26. Apparatus according to claim 24 or claim 25, in which the second state model contains one or more sets.

27. Apparatus according to any of claims 24 to 26, in which the second state model contains two or more leafstates having one or more event driven transitions therebetween.

28. Apparatus as claimed in any of claims 24 to 27, in which one or more of the transitions fires a notification event.

5 29. Apparatus as claimed in any of claims 24 to 28, in which the second state model is implanted over an explicit marker state of the first state model.

10 30. Apparatus as claimed in any of claims 24 to 28, in which the second state model is implanted over an implicit marker state of the first state model.

31. Apparatus as claimed in claim 29 or claim 30, comprising means for deleting the second state model on completion.

15 32. Apparatus as claimed in claim 31, in which the model allows the entering of a state in the first model local to the caller of the second state model only on deletion of the second state model.

20 33. Apparatus as claimed in any of claims 24 to 28, in which the second state model is implanted in free-space.

34. Apparatus as claimed in claim 33, in which the second state model is implanted local to the caller of the second state model.

25 35. Apparatus as claimed in claim 33 or claim 34, comprising means for deleting the second state model on transition to a terminator forming part thereof.

30 36. A method of modelling a state machine substantially as shown in and/or as described with reference of any Figures 10 to 21 of the accompanying drawings.

37. A system for modelling a state machine, the system operating substantially as shown in and/or as described with reference to any of Figures 10 to 21 of the accompanying drawings.

ABSTRACT

STATE MACHINE MODELLING

5 A method of modelling a function call in a state machine comprises
generating a model of a state machine which calls a function call. A
function call mf1 is modelled in a second state machine which is
independent of the first state machine. When the first state machine calls
the function call, for example using a leafstate "calling", the function call
10 state machine mf1 is temporarily implanted over the "calling" state. Static
recursion or infinite compile time recursion is avoided since the implantation
is made only at the time of calling the function call, rather than at compiled
time. After entering the state machine mf1, return is made to the first state
machine after transition to a terminator state, which fires an event called
15 \$return (where \$ indicates scoping back one level, and return indicates an
event which fires the transition to state "after"). This can be described as a
synchronous function call. An asynchronous function call is shown in Figure
17. Here, a second state model is implanted in free-space, and has a
lifetime which is independent of any other state machine model. An
20 asynchronous implanted state machine returns a notification upon
completion, i.e. on transition to a terminator state. Intermediate notifications
may also be given. On exiting any implanted function call state machine,
the implantation is deleted.

25 (Figure 10).




Figure 1

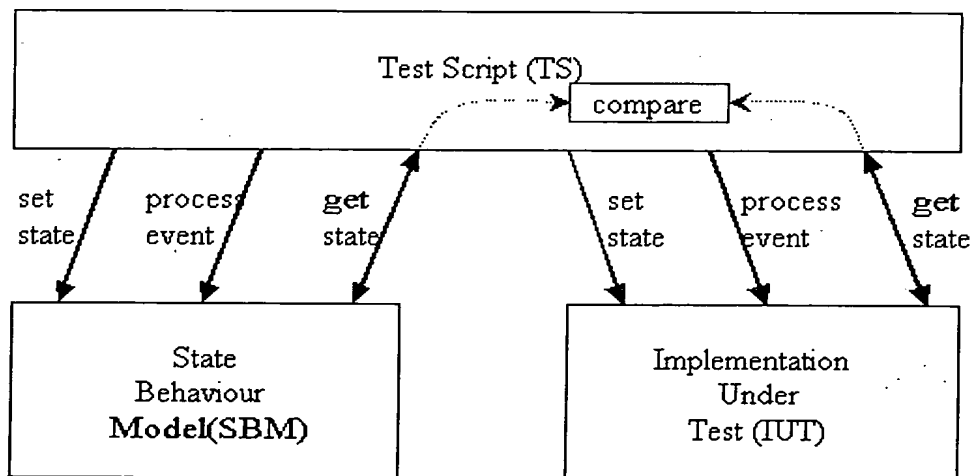


Figure 2A

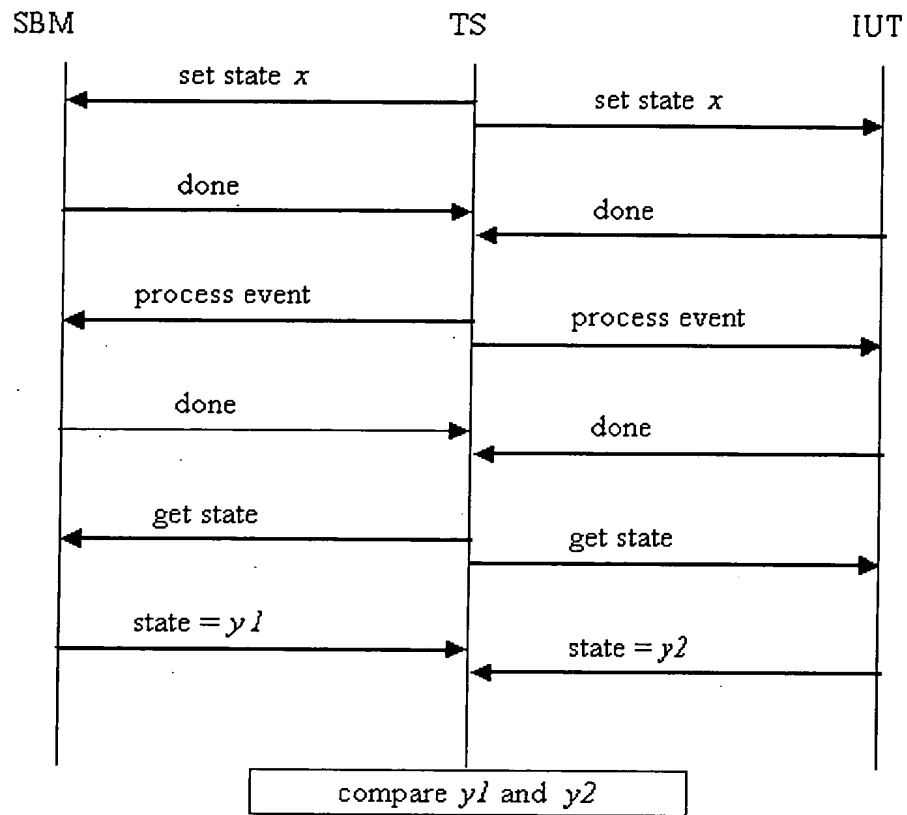


Figure 2B




Figure 3A





Figure 3B



*member is a cluster
(containing two leafstates)
note symbol **a** in the member area*

*alternative : member is a cluster
(containing two leafstates)
note **no** symbol outside the cluster*

*member is a leafstate
note **no** symbol outside the leaf state
(can be useful for self transition
actions)*

*member is a set
(containing two clusters, each of
which contains two leafstates)*

*member is a cluster
(containing a cluster (containing
two leafstates))*

Figure 4

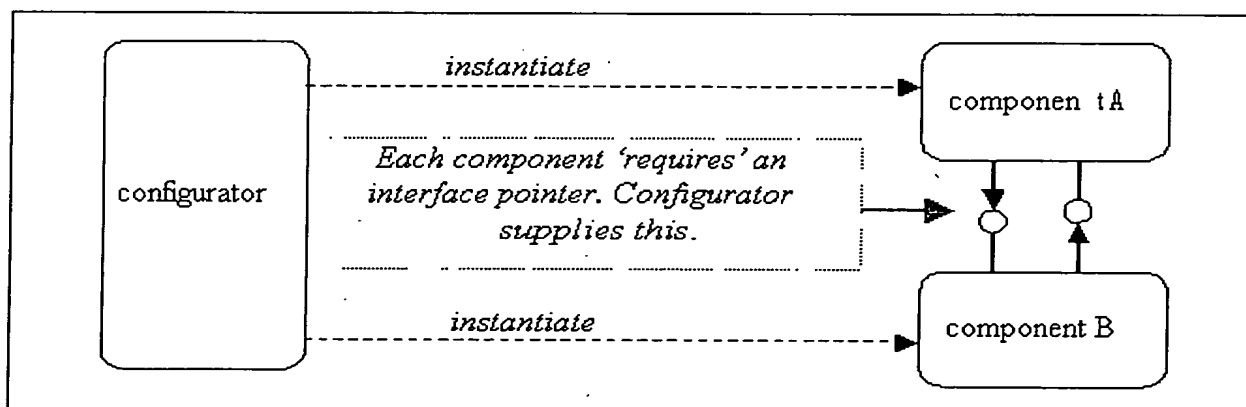


Figure 5

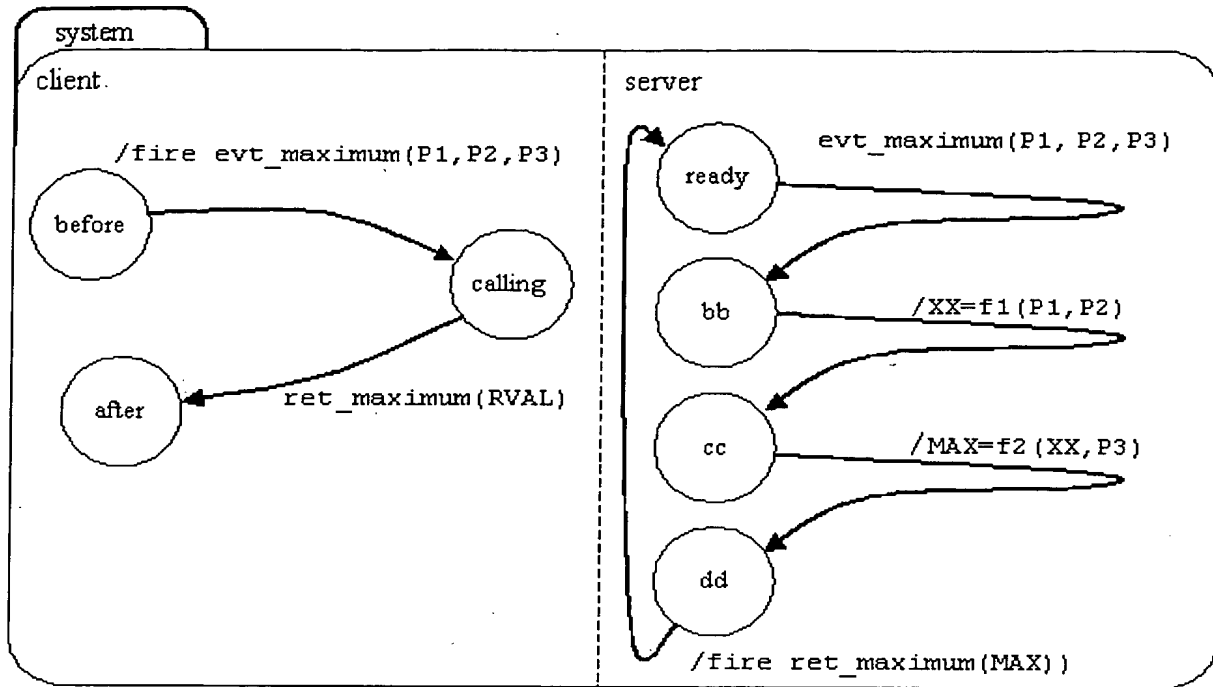


Figure 6

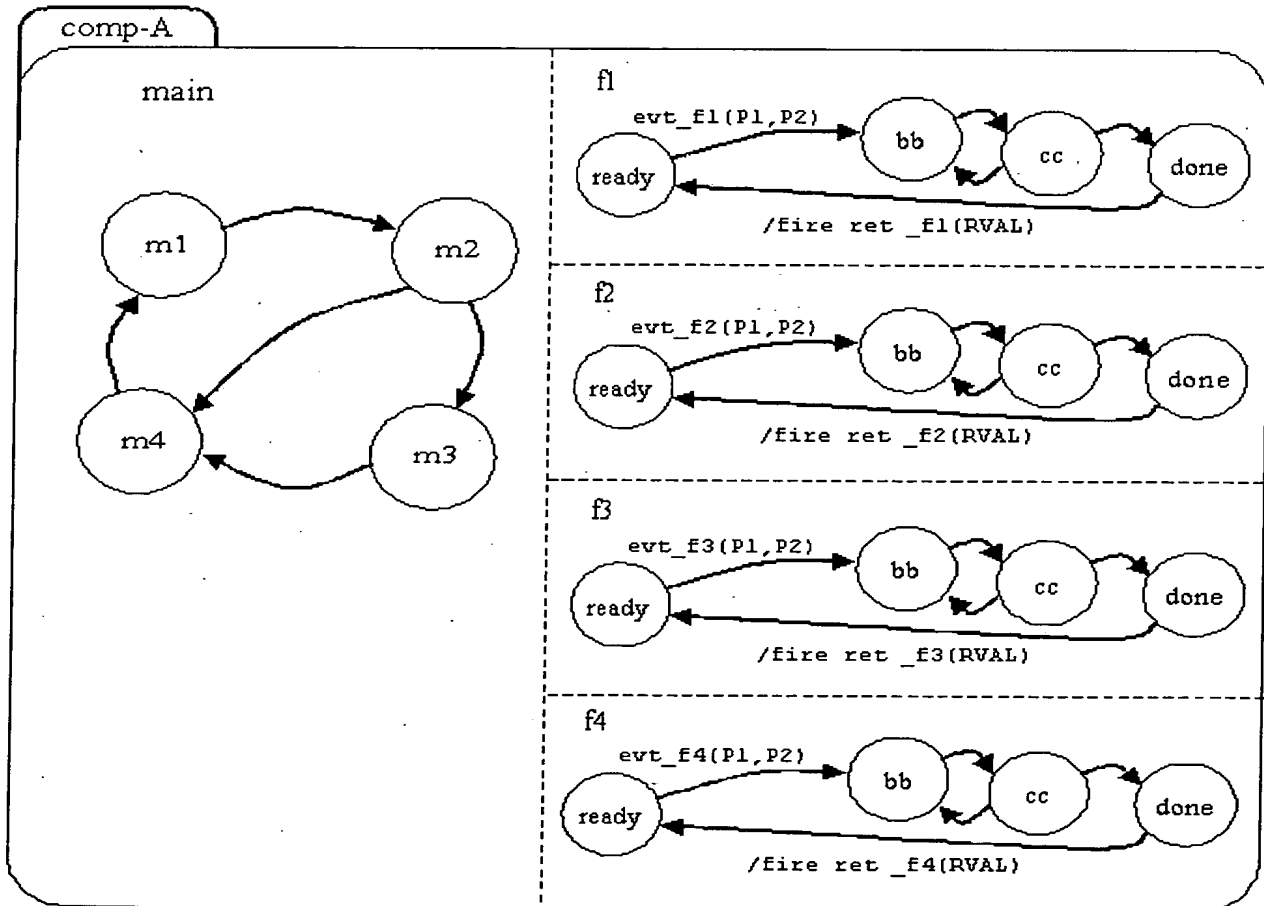


Figure 7

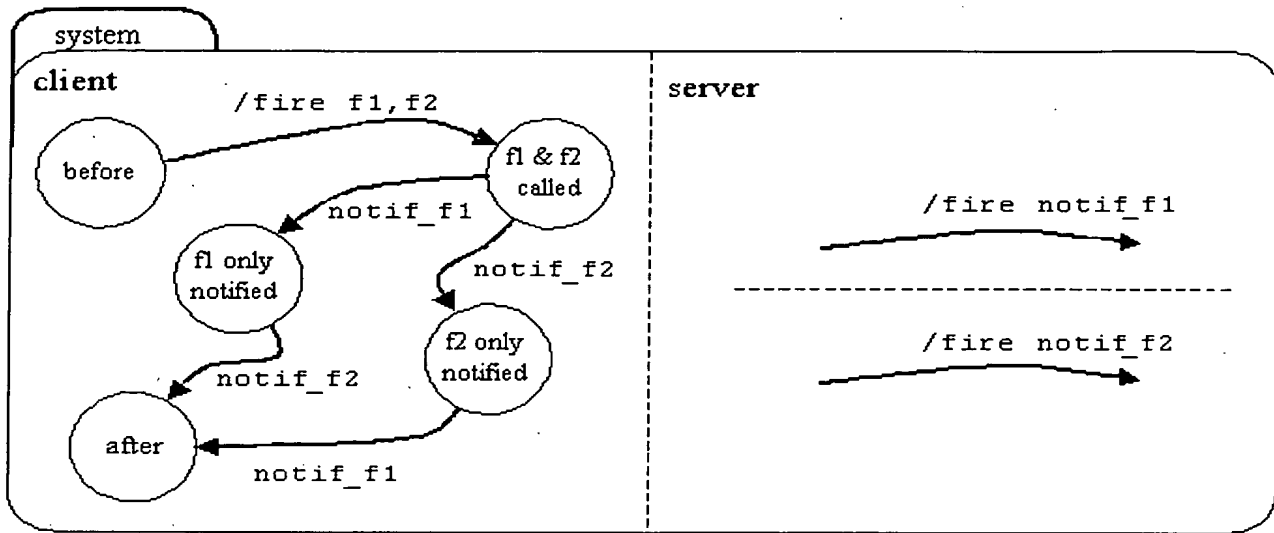


Figure 8

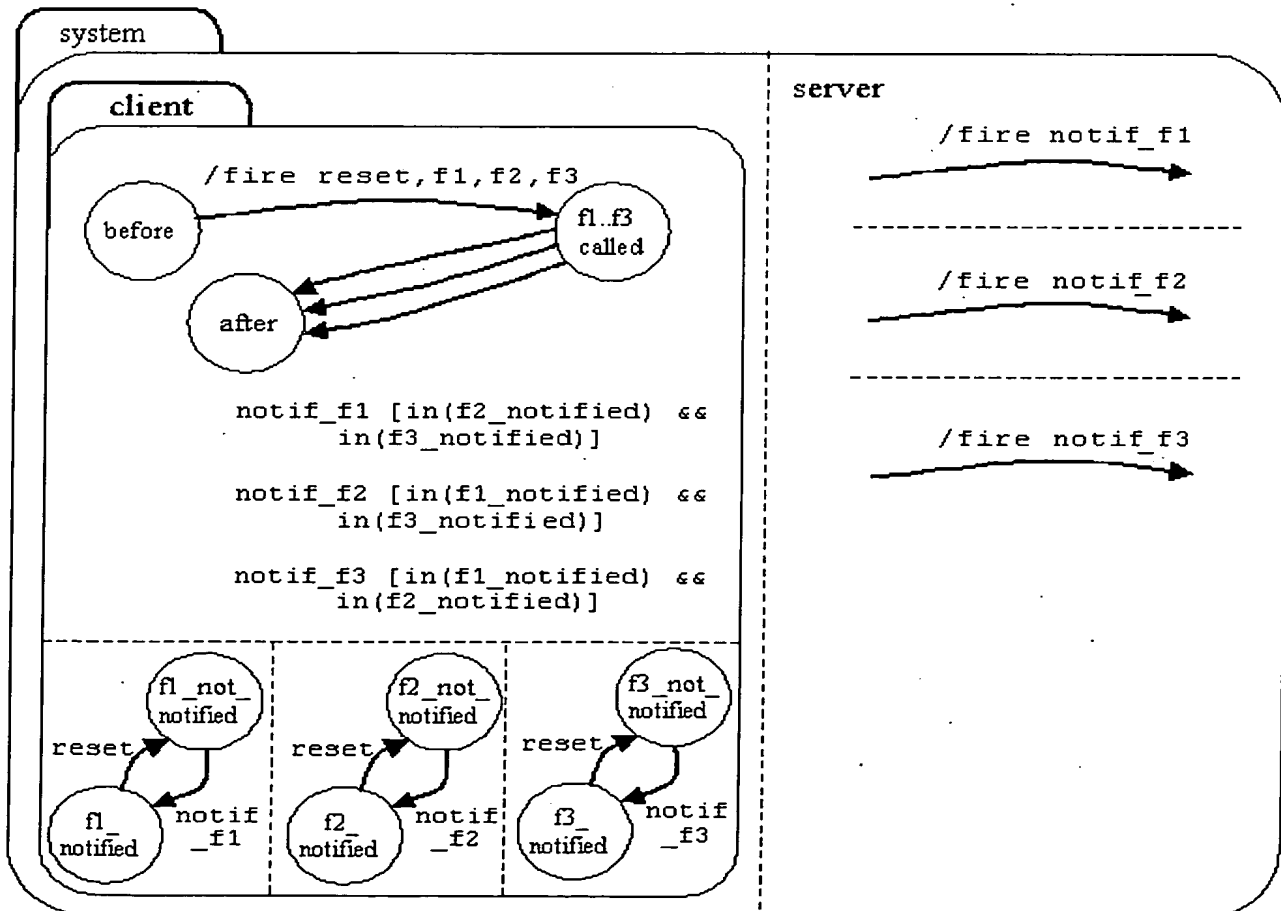
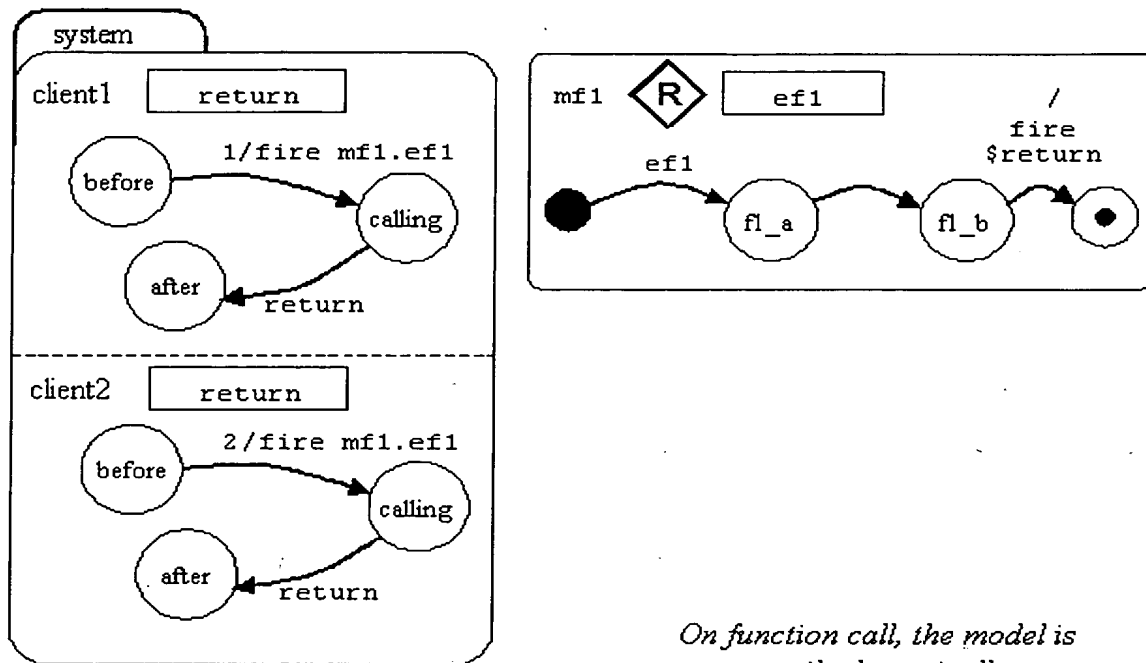


Figure 9



On function call, the model is temporarily dynamically modified into the following:


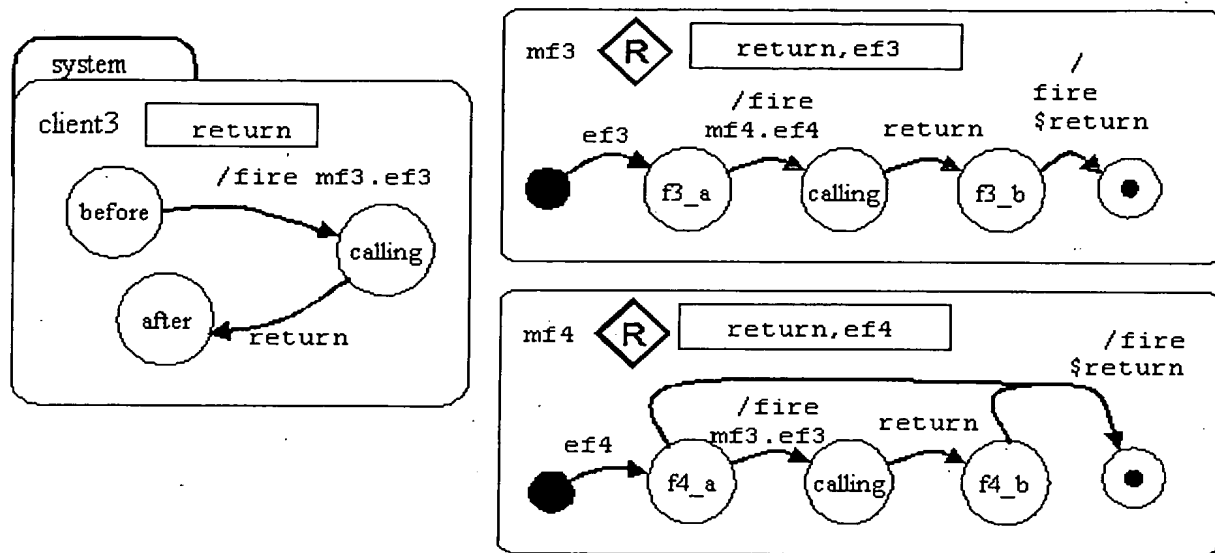


Figure 10



The model is temporarily dynamically modified into the following, at the time when the second invocation of *f3* has taken place




Figure 11

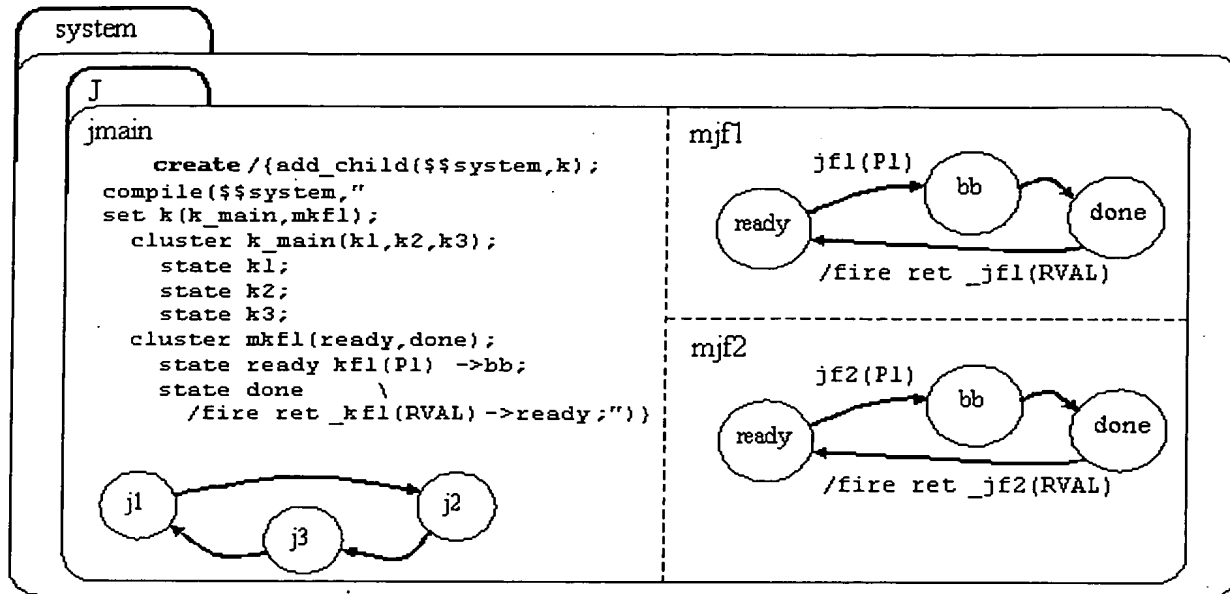


Figure 12

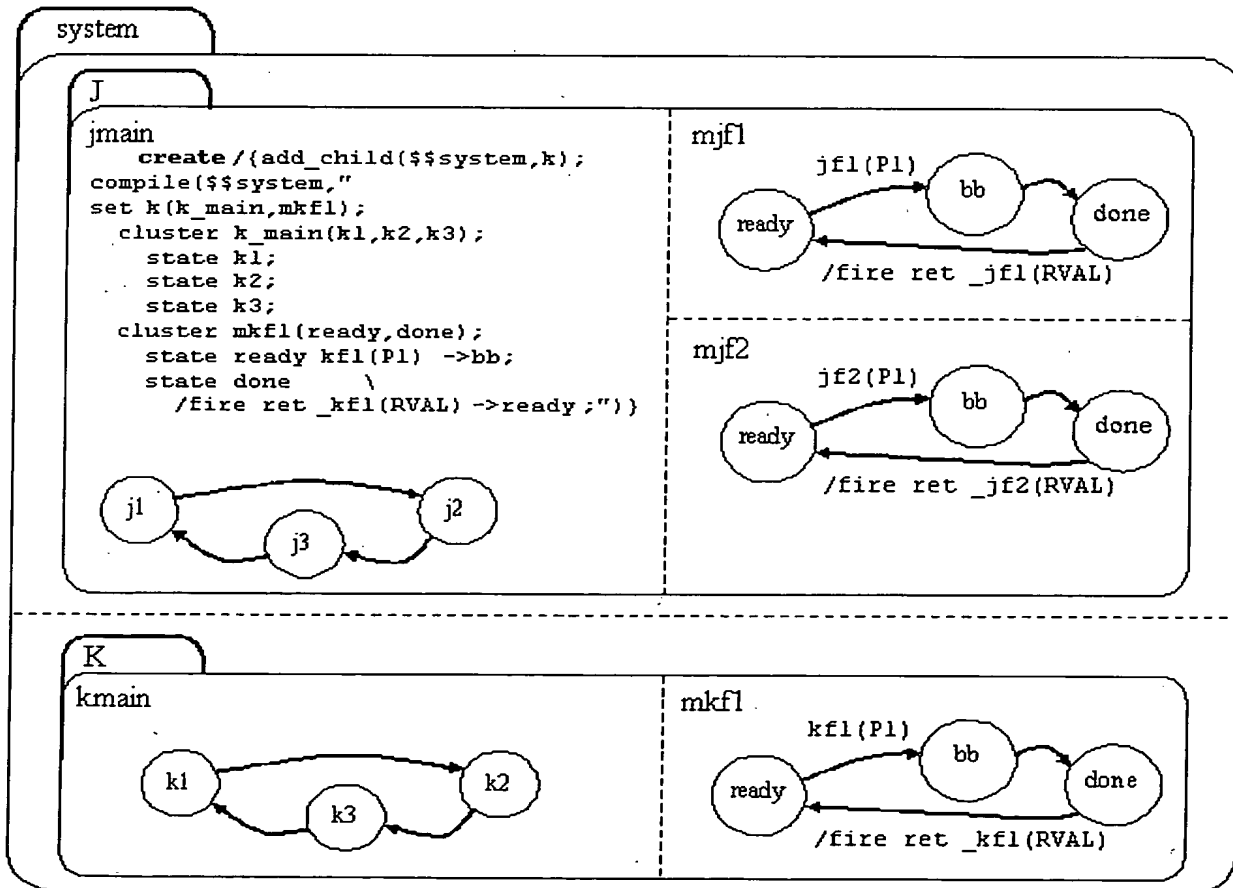


Figure 13

9/13




Figure 14

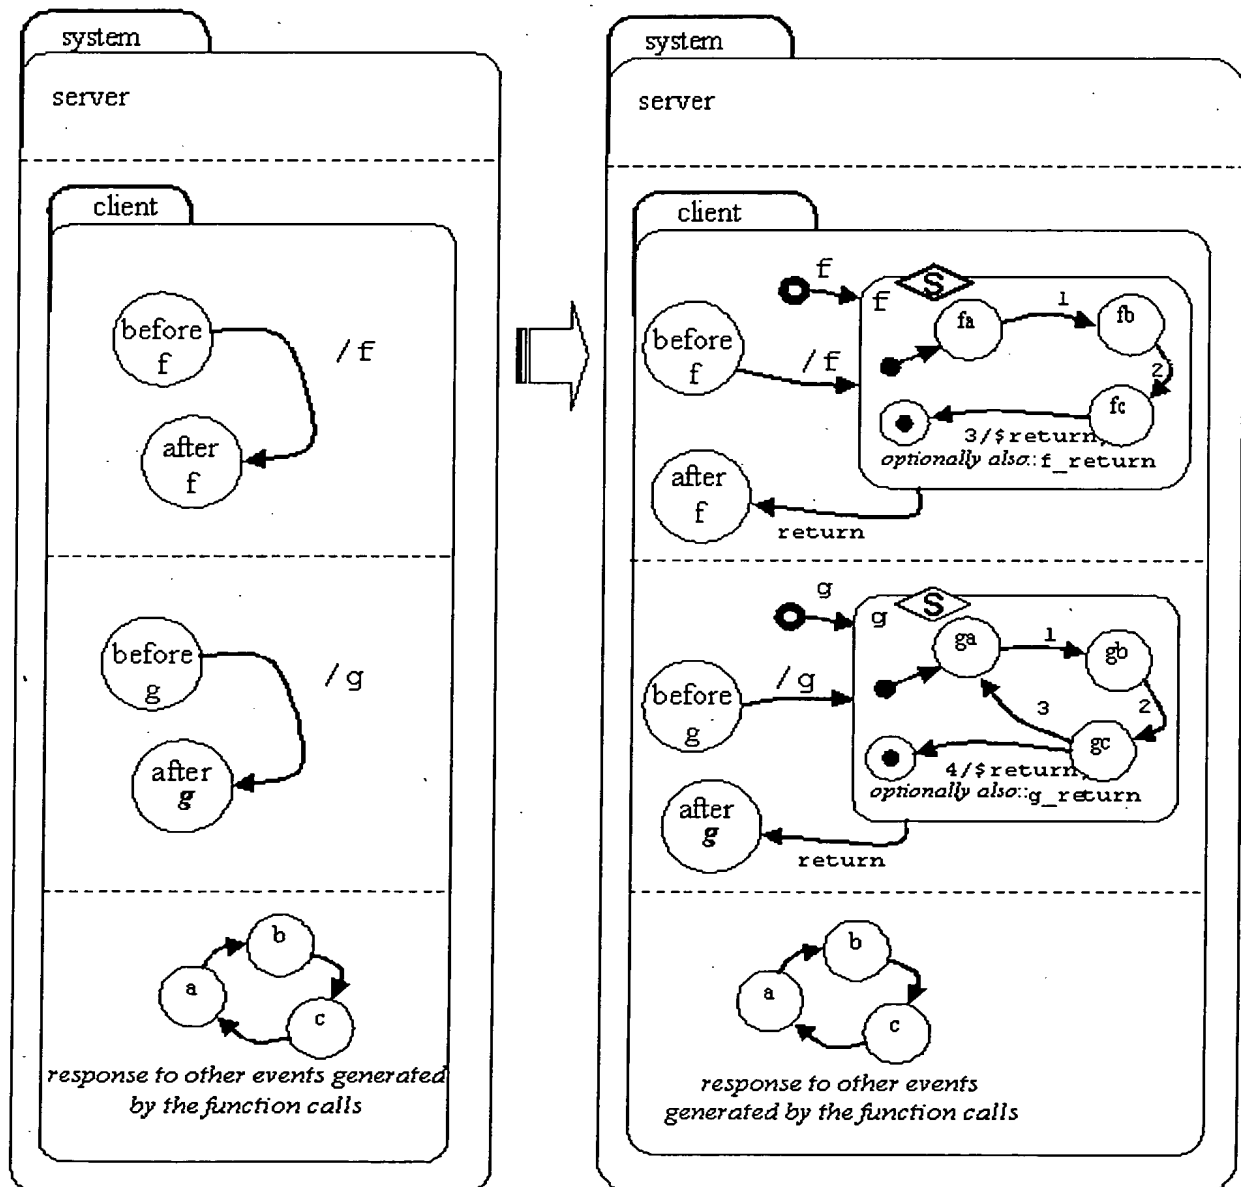


Figure 15

PHGB 020195

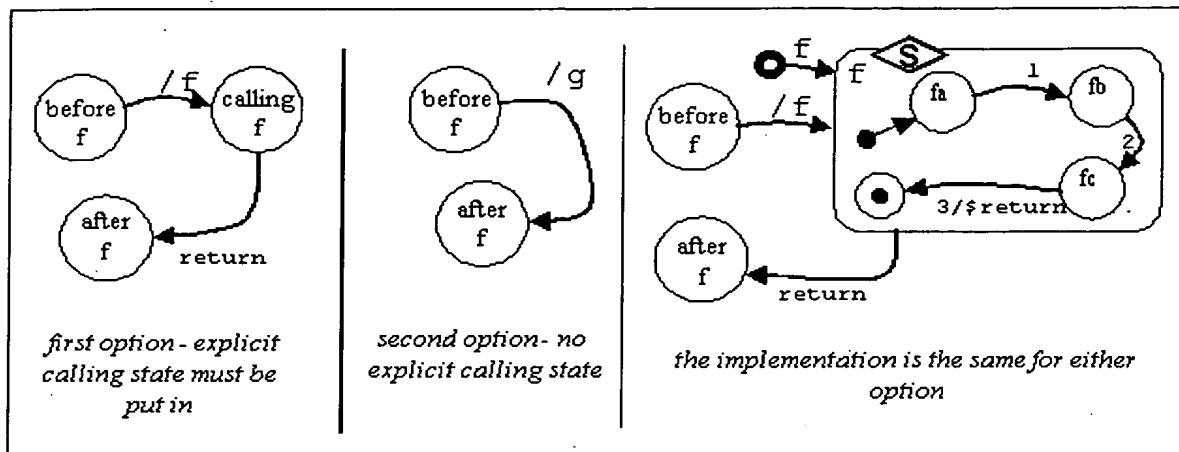


Figure 16

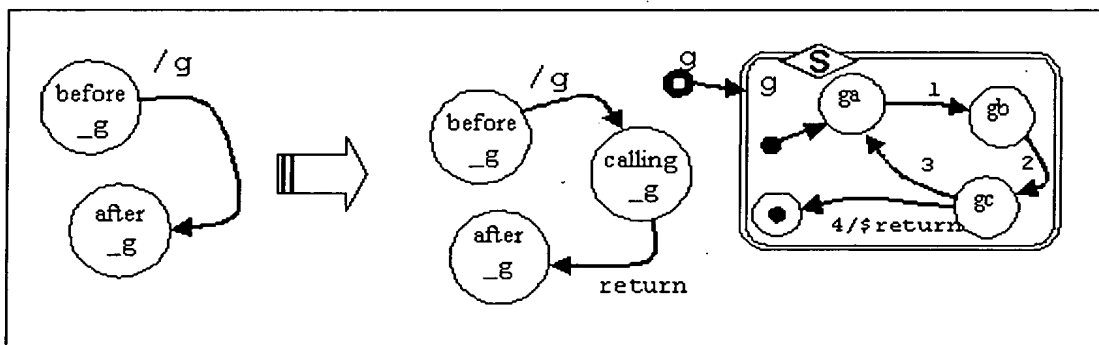


Figure 17




Figure 18

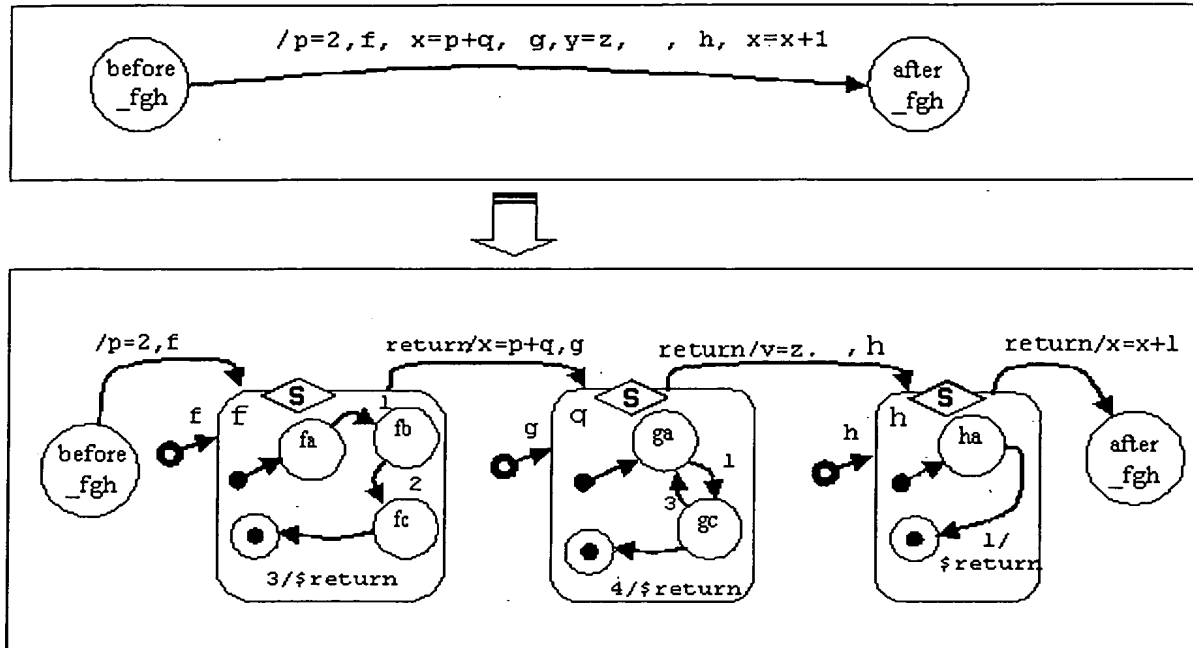


Figure 19

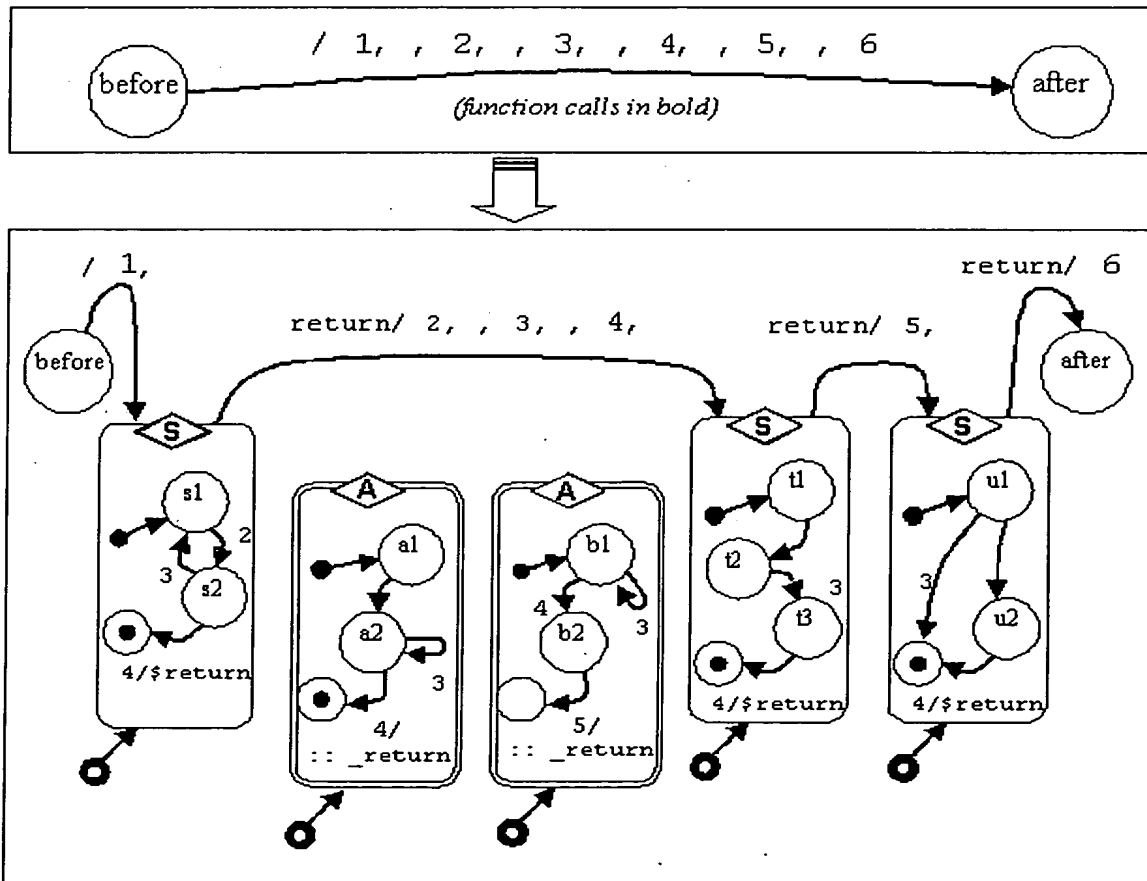


Figure 21




Figure 20




Figure 22

